

Antagonism also Flows through Retweets: The Impact of Out-of-Context Quotes in Opinion Polarization Analysis*

Pedro Calais Guerra, Roberto C.S.N.P. Souza, Renato M. Assunção, Wagner Meira Jr.

Dept. of Computer Science – Universidade Federal de Minas Gerais (UFMG)

{pcalais, nalon, assuncao, meira}@dcc.ufmg.br

Abstract

In this paper, we study the implications of the commonplace assumption that most social media studies make with respect to the nature of message shares (such as retweets) as a predominantly *positive* interaction. By analyzing two large longitudinal Brazilian Twitter datasets containing 5 years of conversations on two polarizing topics – Politics and Sports, we empirically demonstrate that groups holding antagonistic views can actually retweet each other *more often* than they retweet other groups. We show that assuming retweets as endorsement interactions can lead to misleading conclusions with respect to the level of antagonism among social communities, and that this apparent paradox is explained in part by the use of retweets to quote the original content creator out of the message’s original temporal context, for humor and criticism purposes. As a consequence, messages diffused on online media can have their polarity *reversed* over time, what poses challenges for social and computer scientists aiming to classify and track opinion groups on online media. On the other hand, we found that the time users take to retweet a message after it has been originally posted can be a useful signal to infer antagonism in social platforms, and that surges of out-of-context retweets correlate with sentiment drifts triggered by real-world events. We also discuss how such evidences can be embedded in sentiment analysis models.

Introduction

In this paper, we study the implications of the commonplace assumption that most social media studies make with respect to the nature of message shares (such as retweets) as a predominantly *positive* interaction. Given that on general purpose social platforms such as Facebook and Twitter there are no explicit positive and negative signs encoded in the edges, it is commonly assumed (in general, implicitly) that a connection among users through message shares indicate increased homophily among them (Calais et al. 2011; Conover et al. 2011). In general, studies of polarized online communities induced by topics such as Politics and pub-

lic policies do not conduct any explicit analysis of antagonism at the edge granularity, and the degree of separation between communities as well as the controversial nature of the topic is accepted as sufficient evidence of polarization (Garimella et al. 2016). We provide a qualitative and quantitative analysis on the use of retweets as *negative* interactions. In particular, we analyze two large Brazilian Twitter datasets on polarizing topics – Politics and Soccer – which lead us to four main findings related to behavioral patterns on social-media based interactions:

1. Antagonistic communities tend to share each other’s content *more often* than they share content from other less polarizing and conflicting groups. The immediate consequence of this observation is that a simplistic consideration of retweets as an endorsement interaction can lead to misleading conclusions with respect to the nature and polarity of group relationships, as a large number of retweets flowing from one community to another may be misinterpreted as a signal of support.
2. We observe retweets employed as a mechanism for *quoting out of context*, a known strategy of reproducing a passage or quote out of its original context with the intent of distorting its intended meaning (McGlone 2005). In particular, we found that Twitter users share old messages posted by someone from an opposing side with the goal of creating irony when putting the message out of its original temporal context. We observed that some messages are broadcasted even 6 years after they have been originally posted, with the intention of reinforcing an antagonistic and contrary position, rather than indicating support. In our datasets, a significant fraction of retweets crossing antagonistic communities are out of context retweets.
3. As a consequence of Finding 2, messages diffused in a social platform can actually have their polarity *reversed* over time, since the first users sharing the message endorse its original intended content, while other users share the message in response to a real-world event aiming to satirize and to prove that the message’s author was wrong, attaching to it an implicit negative polarity. This concept drift poses interesting challenges for research in text-based sentiment analysis and sarcasm detection.
4. Real-world events can trigger a burst of such out-of-context retweets. We show how the distribution of retweet


response times in a concentrated time span can be a signal which helps detecting sudden sentiment drifts among opinion groups, as they focus on retweeting old tweets from their adversaries during specific real-world events.

We believe the main reason these findings on the use of retweets to convey disagreement remain unnoticed in the social network analysis literature is due the focus on research on bipolarized social networks, characterized by the emergence of *exactly* two dominant conflicting groups, such as republicans versus democrats (Adamic and Glance 2005), pro and anti gun-control (Calais et al. 2013), and pro-life versus pro-choice voices. In this setting, once you determine (automatically or by manual examination) the leaning of a group toward a controversial topic, their (negative) opinion w.r.t. the opposite viewpoint is implicitly determined, and no further analysis of edge polarities is usually performed.


To remove the straight-forward polarity assignment of bipolarized communities and analyze the interplay between retweets and (lack of) antagonism, we collected datasets on discussion domains where more than two communities interact, namely, political discussion in a multipartisan political system and multiple groups of sports fans engaging on conversations about the Brazilian Soccer League. In Figure 1(a), we plot in different colors the three largest communities found in a network of retweets we collected from Twitter during the 2014 Brazilian Presidential Elections, representing groups of people formed around the 3 main candidates (Dilma Rousseff, Aécio Neves and Marina Silva); in Figure 1(b) we do the same for the 12 largest exchanging messages about Brazilian soccer. Differently from bipolarized social graphs, since now there are $K > 2$ possible sides one user may belong to, the identification of an individual as a member of a community does not necessarily imply on antagonism with respect to all the remaining $K - 1$ groups; each group member can be indifferent, or neutral, to a subset of the remaining groups, or even support more than one group simultaneously. As a consequence, we need to conduct a deeper analysis of retweets crossing communities to gain insights on group relationships.

Our work contributes to social media research in two distinct directions. Findings 1 and 3 add to the recent trend on the pitfalls and drawbacks of making inferences based on social media data (Liao, Wai-Tat, and Strohmaier 2016; Rost et al. 2013; Metaxas, Mustafaraj, and Gayo-Avello 2011). Findings 2 and 4, on the other hand, explore how temporal information associated to retweets can be a rich signal to be incorporated into models focused on antagonism detection and real-time tracking of opinions in social media.

In the remainder of this paper, we first discuss related work on polarization and unsigned edges in social networks. Next we analyze two longitudinal Twitter datasets to empirically demonstrate that, on multipolarized social networks, assuming retweets as positive interactions can be misleading. Finally, we characterize how cross-group retweets differ from intra-group retweets with respect to the distribution of the time differences between the message posting time and the retweet action, and we show how this signal can be em-



(a) 2014 Brazilian Political Twitter.



(b) 2010-2016 Brazilian Soccer debate in Twitter.

Figure 1: On the top, a network of retweets obtained from Twitter showing 3 communities formed around the 3 main candidates in the 2014 Brazilian Presidential Elections. On the bottom, communities formed around the 12 top Brazilian Soccer teams. Although both topics are polarizing in nature, in a *multipolarized* domain not every pair of groups is expected to share antagonism.

bedded into models that aim to detect the controversy level among opinion groups and real-time sudden drifts on their sentiment and opinions.

Related Work

On social networks whose edge signs are labeled, antagonistic relationships among communities are naturally reflected by the number of positive and negative edges flowing from the source community to a target community, and the communities themselves can be found by algorithms especially designed to deal with negative edges (Kunegis et al. 2010; Yang, Zhao, and Liu 2015; Lo et al. 2011).

Many works qualitatively discuss and document the empirical observation that unlabeled social interactions on general purpose social platforms such as Twitter and Facebook can convey negative sentiment: replies and comments, as web hyperlinks, do not carry an explicit sentiment label and can be either positive or negative (Leskovec, Huttenlocher, and Kleinberg 2010; Ye et al. 2013). Message broadcasts, on the other hand, have been categorized by early works on behavioral analysis on Twitter as a strictly positive interac-

tion (Boyd, Golder, and Lotan 2010). As users expertise evolved, they had begun finding uses of retweets that do not convey agreement. “Retweets are not endorsements” is a common disclaimer found in biographies of journalists and think tankers in Twitter, whereas some people share stuff that they vehemently disagree only to show the idiocy of the people they oppose. One can also broadcast the original message and append comments to it (“quote RTs”, in Twitter), often in disagreement with the original content, what also contributes to turn shares and retweets into an ambiguous signal with respect to the sentiment they convey (Garimella, Weber, and Choudhury 2016). In summary, retweets and shares are often a “hate-linking” strategy – linking to disagree and criticize, often in an ironic and sarcastic manner, rather than to endorse (Tufekci 2014).

Although documented in the literature as a known behavior, the impact of such “negative” retweets on community and network analysis has not been the focus of in-depth studies so far. Usually, social network analysis practitioners assume, implicitly or explicitly, that retweets (or more generally, shares) have a predominant endorsement nature. A recurrent pattern in community analysis works making sense of social media datasets is that they limit their analysis to social networks whose dominant topic induces a partition of the graph into exactly two conflicting sides: liberal versus conservative parties, pro-gun and anti-gun voices, pro-choice and pro-life (Conover et al. 2011; Livne et al. 2011; Adamic and Glance 2005; Wong et al. 2013). As we will show in the next sections, in bipolarized scenarios, it is harder to grasp the use of retweets to convey disagreement.

Our contribution in this paper is twofold. While we raise awareness to the network science community of the implications of assuming retweets as positive interactions, we propose a new edge-level signal – the retweet response time, i.e. the amount of time the user took to hit the retweet button after the original message has been posted – to help disambiguating positive from negative edges in a social network containing timestamped edges.

Data Collection and Preparation

We used Twitter’s Streaming API¹ to monitor two topics that motivate intense debate on offline and online media and thus are suitable for analysis of formation of antagonistic communities: Politics (Calais et al. 2011) and Sports (Lanagan and Smeaton 2011). Table 1 provides details on the datasets.

In the political topic, our data collection was driven by the main candidates in the 2010 and 2014 Brazilian presidential elections, including Dilma Rousseff, elected for the presidency in both years. In December 2015, Ms. Rousseff faced an impeachment trial conducted by the Brazilian Congress, and on May 12th, 2016, the Senate voted to suspend her for 180 days. The vice-president Michel Temer, elected with her in 2010 and 2014, assumed as the provisory president. We monitored mentions to politician Twitter profiles and names,

Table 1: General description of the two Twitter datasets we consider. Note the large variability on (native) retweet response times.

	Topic	
	Politics	Soccer
period	2010-16	2010-16
# groups	3	12
# tweets	20.5 M	103M
# users	3.1M	8.7M
manual RTs	46K	2K
quote RTs	67K	3K
native RTs	9.1M	30.9M
RT mean response time (hours)	29.5h	43.5h
RT median response time (hours)	0.24h	0.23h
RT response time std (hours)	255.4h	368.7h
# replies	3.2M	20.8 M
reply mean reaction time (hours)	5.1h	3.5h
reply reaction time std (hours)	188.3h	194.0h

the hashtags used by each side participating in the political debate and the names of the presidents of the Brazilian Lower House and the Senate, which directly conducted Ms. Rousseff’s impeachment process in the Congress.

We also collected public tweets about the 2010 to 2016 editions of the Brazilian Soccer League. We monitored mentions to the 12 largest Brazilian soccer teams and match-related keywords, such as “goal”, “penalty” and “yellow card”.

Notice that the fact that we collected tweets during a time span of more than five years allow us to extract the time interval between the original message and each of its retweets, and observe large deltas between these timestamps. We call this time interval the *retweet response time*. Table 1 shows that the mean retweet response time is in the magnitude of several hours and it is an order of magnitude higher than the median retweet response time. Also, its standard deviation is almost an order of magnitude larger than the mean, what indicates a high variability in retweet response times. Compared to replies, the average response time of retweets is about 6 and 12 times higher, in the Politics and Soccer dataset, respectively. This suggests that there might be some specific behavioral and temporal patterns associated with retweets. We will show how such ‘late retweets’ relate to polarization and interactions among antagonistic groups later in this paper.

Three types of retweets can be extracted from the raw JSON tuples: manual retweets, i.e., messages manually created in the format ‘RT @username message’; a quote retweet, when the user prepends or appends a comment to the original message (as in ‘Cool! RT @username message’); and a retweet triggered through the native Twitter retweet button. We have chosen to focus our analysis on native retweets for three reasons:

1. They represent the vast majority of retweets (see Table 1);
2. Although manual and quote retweets are also legitimate

¹Twitter Streaming API:

<https://dev.twitter.com/streaming/overview>.

user interactions, native retweets better reflect how the user interface design affects user behavior, since they are directly implemented in Twitter’s user interface;

3. In a native retweet, the original tweet posting time is provided in the JSON format; therefore we do not need to have collected the original message in order to compute the retweet response time.

Community detection. Once collected we prepared the data for our various analysis as described next.

The first step is to partition the social network induced by the messages and represented as a graph $G(V, E)$ into meaningful communities. Although our methodology does not depend on the specific graph clustering algorithm, finding communities on polarized topics is eased by the fact that it is usually simple to find seeds – users that are previously known to belong to a specific community. In the case of the Twitter datasets we take into consideration, the official profiles of politicians, political parties and soccer clubs are natural seeds that can be fed to a semi-supervised clustering algorithm that expands the seeds to the communities formed around them (Calais et al. 2011; Liao, Wai-Tat, and Strohmaier 2016; Kloumann and Kleinberg 2014).

Different graphs can be built based on the datasets described in Table 1; traditionally, a social network $G(V, E)$ represents a set of users V and a set of edges E that connect two users if they exceed a threshold of interaction activity. The limitation of this modeling is that it hides the individual user-message interactions: for instance, two users holding opposite opinions may propagate different messages from the same media outlet, what could wrongly indicate that both share the same opinion. Connecting users directly hides the fact that the individual messages may have a potentially different sentiment with respect to different entities; i.e., a media outlet may post a positive message w.r.t to a politician one day and a negative message a week later. By representing interactions in a user-message bipartite retweet graph, as shown in Figure 2, we keep this more granular information.

We assume that the number of communities K formed around a topic T is known in advance and it is a parameter of our method. To estimate user and message leanings toward each of the K groups, we employ a label propagation-like strategy based on random walk with restarts (Tong, Faloutsos, and Pan 2008): a random walker departs from each seed and travels in the user-message retweet bipartite graph by randomly choosing an edge to decide which node it should go next. With a probability $(1 - \alpha) = 0.85$, the random walker restarts the random walking process from its original seed. As a consequence, the random walker tends to spend more time inside the cluster its seed belongs to (Calais et al. 2011). Each node is then assigned to its closest seed (i.e., community), as shown in the node colors in the toy example from Figure 2. For more details on the random walk-based community detection algorithm, please refer to (Calais et al. 2011).

For both the Politics and Soccer Twitter dataset, we performed a validation of the K communities we found using a sampling strategy on the correlation between communi-




Figure 2: A bipartite user-message graph connecting users with messages they interact with. To find communities, we run a random walk with restarts from each seed that represents a community (notice in the figure that they are explicitly labeled); the random walker will traverse more frequently the links and nodes belonging to the community the seed belongs to. Node colors represent relative proximities to the red/blue sides.

ties and profiles that make explicit their side. In particular, we exploit the evidence provided by many Twitter users that append to their profile names the soccer team or political party they support; and, in general, the content they publish will favor the respective mentioned side, as we observed through manual inspection of a sample. For example, @[first name][name of favorite soccer team] is a common account name pattern, as in @JohnCruzeiro, through which John declares he is a Cruzeiro fan. From the 13,892 messages these specific users generated in the Elections and Soccer dataset, we found that 91.53% were assigned by our algorithm to the community indicated in their profile name. Although we acknowledge that user account names are an imperfect ground truth and these users are more likely to present a more active and clearly defined behavior and thus they are more easily classified, we believe this number indicates that the accuracy of the random-walk based clustering method is enough for our data analysis purposes.

Finding 1: antagonistic groups retweet each other more than they retweet other groups

As we pointed out in Section 1, the polarity relationships among the K communities found is not an explicit byproduct of a community detection method whose input is an unsigned graph. Recall that, on bipolarized domains, no subsequent analysis is usually performed, other than the quantification of the degree of separation between the pair of communities, using commu-

nity quality metrics such as *modularity* (Livne et al. 2011; Adamic and Glance 2005). It is a standard practice to assume that the more separated the communities are, the more antagonism is observed, as a consequence of the homophily principle (McPherson, Smith-Lovin, and Cook 2001).

The intrinsic limitation of a bipolarized network is that only one separation metric value can be computed, since there is only one pair of communities. Since we are studying $K > 2$ cases, we now have $\binom{K}{2}$ pairwise community metrics to compare. For the sake of simplicity, for each pair of communities we compute the proportion of retweets triggered from users belonging to community i that flow toward messages posted by members of community j relative to all retweets that community i trigger to the other groups in the graph:

$$RT_ratio(i, j) = \frac{RT_{i,j}}{\sum_{k=1, k \neq i}^K RT_{i,k}} \quad (1)$$

We compare $RT_ratio(i, j)$ considering the known local rivalries that exist in Brazilian Soccer among soccer clubs from the same Brazilian state, as listed in Table 2.

Table 2: Local rivalries in Brazilian Soccer. Stronger antagonism exists between soccer clubs and communities of supporters belonging to the same Brazilian state.

Brazilian state	local rivalries
Minas Gerais	Cruzeiro, Atlético
São Paulo	SPFC, Santos, Corinthians, Palmeiras
Rio G. do Sul	Grêmio, Internacional
Rio de Janeiro	Flamengo, Fluminense, Vasco, Botafogo

In Figure 3 we plot $RT_ratio(i, j)$ for all the $\binom{K}{2}$ pairs of communities formed around supporters of Brazilian soccer clubs, and we visually discriminate between pairs of rival communities (red triangles) and non-rival communities (green circles) according to the ground truth from Table 2. The graph shows a somewhat unexpected result: pairs of communities that are *more antagonistic* (i.e., the opposing sides belong to the same Brazilian state) tend to retweet each other’s content *more often* than when there is less, or no antagonism between them. For example, Cruzeiro’s community (id = 8) targets about 65% of its cross-group retweets to Atlético’s community, their sole fierce rival in Brazilian state of Minas Gerais. As another example, community 1, which identifies supporters from Rio de Janeiro team Flamengo, prefers to retweet messages for their three local rivals. As a general rule, red triangles dominate green circles, i.e., retweets are targeted more often to antagonistic communities than to more neutral, less conflicting groups.

The fundamental insight to learn from Figure 3 is that retweets carrying a negative polarity directly impact the network structure and make antagonistic communities *closer* in the social graph. On traditional bipolarized domains in which current literature focuses, this apparent paradox is inherently unnoticeable, since there is only a single pair of




Figure 3: $RT_ratio(i, j)$ for each pair of 12 communities discussing Brazilian soccer in Twitter. More antagonistic communities retweet each other more than neutral, less polarizing communities.

antagonistic communities and thus only a single separation metric to be computed.

We list a few intents that motivate Twitter users in retweeting messages they disagree with:

- **Share to show contrary opinion.** Many times, a user propagates a message he or she disagrees with to show the message to their followers or friends and comment on that content. The goal is to start a discussion and gauge reactions.
- **Fake or edited retweets.** We do not include these retweets in our analysis, but some Twitter users create fake retweets, in the format “RT @user fake message”, assigning to @user a message that has never has been posted. Fake retweets have already being investigated as a spamming activity in Twitter (Mowbray 2010), in which spammers try to borrow from the reputation of celebrities. In the context of polarized discussions, however, the goal is different – to make criticism or even spread false information (Mustafaraj and Metaxas 2011).
- **Out-of-context quoting.** We will provide an in-depth analysis of this behavior in the next section. In summary, a user propagates a message he or she disagrees with and puts it out of context, in order to create sarcasm or irony. In this case, we usually see messages being shared long after they were originally posted, typically when the original message stated a prediction that turned out to be false later.

Negative retweets and the filter bubble. In a recent study by Pew Research Center, polarized discussions have been identified as one of the top 6 most common conversational structures in Twitter (Smith et al. 2014). For that reason, better understanding the social structures induced by polarized debate is important because polarization of opinions induces segregation in the society, causing people with different viewpoints to become isolated in islands where everyone thinks like them (Vydiswaran et al. 2012). Such filter

bubble caused by social media systems limits the exposure of users to ideologically diverse content, and is a growing concern (Lazer 2015; Bakshy, Messing, and Adamic 2015). The behavioral pattern we document here has the unintentional side effect of *reducing* the filter bubble, letting followers of advocates of one viewpoint to get to know the opinions of the other side.

The “paradox” of antagonistic communities being linked by more retweets make clear some assumptions which are commonly implicitly made in the literature with respect to the treatment of edge signs. While the correctness and applicability of each assumption depends on inherent characteristics of each dataset, we advocate that it is a good practice to make it clear the expectations with respect to the following aspects/metrics:

1. **Edge sign prior.** The vast majority of community detection methods on social media networks are built over the assumption – which is, most of the time, not make explicit – that there is an apriori knowledge that edges are more likely to be positive than negative. If $P(\text{sign}(\text{edge}) = +)$ is sufficiently high, it is reasonable to expect that the method will output the identification of groups of users and messages around a cohesive viewpoint and high level of homophily. For instance, in a blog citation network, one blog may cite the other to disagree with it, but since most of the time a blog citation is an endorsement rather than a disapproval, edge label-agnostic community detection methods work reasonably.
2. **Antagonism and community separation metrics.** It is a standard practice to measure the degree of antagonism between communities through separation metrics such as modularity, considering that the more separated the communities are, the higher their level of antagonism and controversy (Adamic and Glance 2005; Conover et al. 2011). However, a smaller modularity may actually indicate an increase of antagonism through interaction via negative retweets and debate through replies and comments.
3. **Domain of discussion and antagonism.** The other implicit assumption usually made by social network analysis researches on networks subject to polarization is that the domain implicitly denotes antagonism, rather than being inferred from a principled method that analyzes the network structure and content. More formally, it can be assumed that, once you condition on edges that cross communities, the likelihood of an edge being negative is now greater than being positive. As a consequence, once users are grouped into two communities, members of one group will automatically be assigned to have a contrary or antagonistic opinion regarding the remaining group. These works do not deal with differences between antagonism or indifference, neither with a more accurate handling of edge signs.

In the next section, we will use the temporal context where retweets occur as evidence that indicates which retweets have a higher probability of conveying antagonism.

Finding 2: out-of-context retweets are more prevalent on cross-group relationships

We are now interested in understanding differences between internal retweets, i.e., those which connect users and messages belonging to a single community, and cross-group retweets, i.e., those which are triggered by users from one community but propagate a message posted by an user from another group.

We focus our analysis on the retweet response time – the time interval between the original message posting time and the retweet time. Previous studies found that 50% of retweets tend to occur up to one hour after the original message posting time (Kwak et al. 2010); other studies have related very short and very long retweet response times to fraudulent activity to boost user popularity (Giatsoglou et al. 2015). Our goal is to analyze retweet response time under the perspective of the message polarity and the polarity that the user broadcasting the message is attempting to convey.

In Figure 4 we plot the cumulative distribution of retweet response times, measured in seconds. We plot this distribution for internal (intra-community) and cross-group (inter-community) retweets for both the Soccer and Politics dataset. Notice that cross-group retweets tend to occur later when compared to internal retweets. For instance, at least 30% of retweets connecting groups in both datasets occur after 16 hours of the original message posting time; on the other hand, in the case of internal retweets, only 10% of retweets occur temporally far from the original post. Notice, also, that the four curves group into two clusters, indicating that in both topics the retweet response time distribution is similar.




Figure 4: On average, retweets which cross antagonistic communities tend have larger response times than inter-community retweets. This empirical observation suggest the potential use of retweet response times as a qualifying signal for prediction of edge labels and community memberships.

We now take a closer look at some messages. For instance, consider the following tweet posted by the official account of the Brazilian elected vice-president Michel Temer about

Table 3: The top 5 most retweeted messages from Brazilian VP (@MichelTemer) during impeachment voting period were very old retweets. Users retweeted old messages indicating support from Temer to Dilma, although the moment was of tension and conflict between them.

tweet	# retweets	avg. retweet response time (days)
<i>"We will shout loud to everyone: "Dilma is our President"".</i>	9,669	606
<i>"Impeachment is unthinkable and has no basis in law neither in Politics."</i>	9,338	385
<i>"Dilma is the best person to conduct our country."</i>	5,031	628
<i>"Congratulations on your birthday, Dilma. God Bless You."</i>	2,020	857
<i>"Dilma is displaying confidence and knowledge."</i>	1,627	2105

Table 4: 2 of the top 5 most retweeted tweets from Brazilian President (@dilmabr) during impeachment voting period were very old retweets, indicating support from Dilma to Temer. Dilma, however, were accusing her VP to plan a coup against her.

tweet	# retweets	avg. retweet response time (days)
<i>"I thank my VP Michel Temer for all the support."</i>	4,314	538
<i>"The impeachment is against the wishes of the Brazilian people."</i>	3,635	1.21
<i>"Follow President Dilma live from Periscope."</i>	684	0.58
<i>"President Dilma will make a speech on the Brazilian Senate decision."</i>	606	0.39
<i>"Our VP @MichelTemer is now on Twitter. Let's welcome him!"</i>	329	693

a speech given on TV by his presidential candidate, Dilma Rouseff, during the 2010 Presidential Elections:

2010-08-05 11:11 PM: @MichelTemer: *Dilma is displaying confidence and knowledge.*

Six years after this post, President Rouseff has been suspended by the Brazilian Congress following an impeachment trial of misuse of public money. In response, she gave a speech on March 12th, 2015 accusing VP Temer's party (PMDB) to plan a coup against her. During her speech, many users contrary to Rouseff began retweeting Temer's 2010 message:

2016-05-12 12:23 AM: @randomRouseffOppositor: RT @MichelTemer: *Dilma is displaying confidence and knowledge.*

This is a clear attempt to retweet a message with the intention to attach to it a negative connotation; it does not support nor endorse its original content. On the contrary, retweeters of this message in 2016 attach to it a semantics which is exactly the opposite to the one stated in the direct interpretation of the message, what is precisely the definition of irony (Wallace 2013). While the "contextomy" practice usually refers to selecting specific words from their original linguistic context (McGlone 2005), we see that, in Twitter, such change of meaning is usually associated with some temporal evolution.

Politicians are often targeted by out of context quotes (Boller and George 1989). Tables 3 and 4 list the most popular tweets from @MichelTemer and @dilmabr which received retweets during the impeachment voting process period. In case of VP Temer, all top 5 most retweeted tweets are very old tweets; and the same applies to 2 of the top 5 messages from Dilma Rouseff. All those messages indicate affective and positive relationships among both politicians, even though the moment was of conflict between them due to the impeachment trial. As a consequence, content-based and network-based algorithms built over the retweet-

as-endorsement assumption can easily be led to make wrong predictions over this data.

Late retweets and Twitter user attributes

To further explore how retweet response times can be an explanatory signal that helps on various social-related prediction tasks, we investigate how late retweets are disproportionately targeted to some types of Twitter users. In particular, we calculated the prevalence of late retweets targeting messages posted by three types of users:

1. Verified users; i.e., users who own a blue verified badge assigned by Twitter to let people know that an account of public interest is authentic. In the Politics dataset, only 17% of the retweets target verified users.
2. Users who have a large follower base; we classified in this category users who have at least 100,000 followers. In the Politics dataset, 23% of retweets target such users.
3. Users who have been retweeted by users who were also retweeted by them. In the Political dataset, only 2% of retweets are triggered by reciprocal retweeters.

For the sake of this analysis, we considered a retweet to be "late" if its response time is at least two standard deviations greater than the average response time. In Figure 5, we observe that, when compared to "early" retweets, late retweets disproportionately target messages from verified users, and users who have a large follower base. In both cases, more than two thirds of late retweets target those types of users. Furthermore, we see that users who mutually retweet each other are less likely to be targeted by a late retweet.

Those measures reinforce a few hypotheses. The first is that late retweets are most commonly targeted to famous and well-known users because they provide *context* to support the ironic and sarcastic purpose of retweeting their tweets out of their original temporal context. Second, the observation that mutually-retweeted users are less likely to be involved in a late retweet is an indication that late




Figure 5: Late retweets are disproportionately targeted to users owning verified accounts and a large follower base. 67% and 77% of late retweets target verified and large-follower based users, respectively. On the other hand, reciprocal retweeters are less often involved in late retweets. Results are similar in the Soccer dataset.

retweets tend to be negative interactions, since reciprocal interactions have been shown to be correlated to homophilic ties (Weng et al. 2010; Kwak et al. 2010).

While in isolation it is hard to tell whether a retweet is an endorsement or not, new signals captured from the social and temporal context, such as the retweet response time, can help on the design of community detection methods. Negative retweets also pose challenges to signed network analysis: as we showed, the sign of an edge in the social graph may actually depend on *when* the edge has been created, what suggests that embedding temporal information on edge creation information may enhance signed network models and algorithms that focus on prediction tasks such as social tie predictions.

Finding 3: Message polarities may reverse over time

One implication of out-of-context retweets is that messages' polarities can actually *reverse* over time. Consider this tweet posted by a popular profile representing the Brazilian soccer club Atlético Mineiro posted in early 2013 mentioning their rivals Cruzeiro:

2013-02-02 10:20 PM @caatleticomg: 2013 will be a great year for Cruzeiro: financial debt and injured players.

“Great”, here, was employed in an ironic way: the tweet was actually predicting (and wishing) a bad year for its rival Cruzeiro. At that year, however, Cruzeiro enjoyed one of the best league performances of its history, winning the national league after scoring 76 points, eleven more than the runner-up. In Figure 6, we show the proportion of retweets of this message originating from Cruzeiro’s supporters over time; the original message was posted at time 0. Notice that 400,000 seconds (277 days) after the message has been originally posted, there is a sudden drift in the ratio of retweets

originating from Cruzeiro’s supporters; it goes from a negligible ratio to about 95% of retweets. The change on the dominant group retweeting the message happened when Cruzeiro won the Brazilian National League and fans were celebrating, and they wanted to make clear that the ironic prediction from its rivals have flagrantly failed.

Since the same content can convey an opposite sentiment depending of its temporal context, text-only irony and sarcasm classifiers such as (Joshi et al. 2016) will not be able to correctly predict the intent of the message propagator. In fact, context plays a significant role on human communication (Wallace 2013) and the polarity reversal we witness here calls for more context-aware signals on sarcasm detectors, what includes temporal and social features in models.




Figure 6: Message polarity may reverse over time: a message initially negative to Cruzeiro’s supporters has turned into positive, after being retweeted with an ironic intention after Atlético fans predictions on Cruzeiro performance have failed.

Finding 4: Spikes of late retweets correlate with sentiment drifts

In Section 5, we showed that out-of-context retweets have an increased chance of being a negative interaction. We now investigate whether there is a concentration of such retweets in specific time frames. We focus on the Soccer topic, more specifically, in the year of 2013, which was particularly eventful for Atlético and Cruzeiro supporters.

We group messages at a daily granularity and its source (Atlético or Cruzeiro fans). For each of these sets, we plot in Figure 7 the 95th percentile of the retweet response times of the messages posted by each group on that day. We notice that the main events related to the Brazilian soccer world were captured as spikes: in July 2013, Atlético won its first Copa Libertadores, what generated a huge of spike of retweets of Cruzeiro supporters who tweeted that Atlético would never win the competition.

The remainder of the year was not favorable to Atlético, though. In November, Cruzeiro won the national league, and in December Atlético lost the FIFA Club World Cup. The sequence of unfortunate events for Atlético fans coincides




Figure 7: 95th percentile of retweet response times triggered by each day during 2013. Spikes coincide with significant real-world events that triggered different reaction on antagonistic groups; in general, we observe that wrong predictions made by rival communities are retweeted by rivals when they are proved wrong.

with a series of spikes in retweets of their old messages by Cruzeiro fans; including the wrong prediction that 2013 would be a “great” (ironically meaning “bad”) year for the Cruzeiro club.

As on Finding 3, we also see potential for using the temporal information associated to retweets to enrich real-time sentiment analysis models, and we leave a more thorough exploring of retweets response times in sentiment analysis algorithms as future work.

Conclusions

In this paper we explore the observation that, in the vast majority of social media studies, especially those based on Facebook and Twitter data, there is no explicit positive and negative signs encoded in the edges. Since inferring individual edge polarities in a unsigned graph is not a trivial task, most social studies assume that retweets and shares are endorsement interactions. No specific analysis on the polarity of the links crossing the communities is usually conducted and antagonism is assumed due to the modular division of the social graphs into two communities historically known to be antagonistic, such as democrats and republicans.

Although very recent papers on retweeting activity still qualify retweets as a strictly positive interaction (Garimella et al. 2017; Metaxas et al. 2015; Liu and Weber 2014), we show that retweets can actually carry a negative polarity, conveying a sentiment which is opposite to the one explicated in the tweet’s text. We believe the neglected impact of negative retweets explain, in part, the low accuracy levels obtained in some user polarity classification experiments (Cohen and Ruths 2013). We also demonstrate that negative retweets contribute to make antagonistic groups *closer* to each other in a network of retweets, what can lead to misleading conclusions by naïve network models, particularly when multiple communities

are found, due to the absence of retweets between neutral communities.

We found that one of the reasons that motivate Twitter users to broadcast tweets they disagree with is to create irony by broadcasting a message in a different temporal context, especially when a real-world event that disproves the original message argument happens. Such behavior finds similarity on *quoting out of context*, a practice already described in the Communications literature (Boller and George 1989).

We believe the better understanding of retweets as multifaceted social interactions which can be (1) possibly negative and (2) have a temporal component may support the design of algorithms that exploit the network structure in conjunction with opinionated content to better perform tasks typically offered by social media platforms, such as content recommendation, event detection, sentiment analysis and news curation (Calais et al. 2011; Tan et al. 2011).

We acknowledge that one of the limitations of our study is that the method that find clusters through random walks from seed nodes do not distinguish between positive and negative retweets; then, some users may be wrongly classified exactly due to the ironic broadcasts he may engage in. However, since positive retweets are still dominant, this effect should affect a few users. Nevertheless, we can think of algorithms that simultaneously infer both edge polarities and user memberships as an interesting future work. Another interesting approach would be weighting edges by their retweet response times; community detection methods could give more priority to recent retweets when seeking for homophilic relationships.

Our work also reinforces the opportunity and possibilities of building rich models which combine content, network structure and temporal dimensions of the underlying social data. Since each dimension is ambiguous in nature, powerful predictive and descriptive methods can be built upon combining these three evidences.

Acknowledgments

This work was supported by CNPQ, Fapemig, InWeb, MASWeb, BIGSEA and INCT-MCS.

References

- [Adamic and Glance 2005] Adamic, L. A., and Glance, N. 2005. The political blogosphere and the 2004 u.s. election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*, LinkKDD ’05, 36–43. New York, NY, USA: ACM.
- [Bakshy, Messing, and Adamic 2015] Bakshy, E.; Messing, S.; and Adamic, L. 2015. Exposure to ideologically diverse news and opinion on facebook. *Science*.
- [Boller and George 1989] Boller, P. F., and George, J. H. 1989. *They never said it : a book of fake quotes, misquotes, and misleading attributions*. Oxford University Press New York.
- [Boyd, Golder, and Lotan 2010] Boyd, D.; Golder, S.; and Lotan, G. 2010. Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. In *Proceedings of the 43rd Hawaii International Conference on Social Systems (HICSS)*. IEEE.

- [Calais et al. 2011] Calais, P. H.; Veloso, A.; Meira, Jr, W.; and Almeida, V. 2011. From bias to opinion: A transfer-learning approach to real-time sentiment analysis. In *Proc. of the 17th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.
- [Calais et al. 2013] Calais, P. H.; Jr., W. M.; Cardie, C.; and Kleinberg, R. 2013. A measure of polarization on social media networks based on community boundaries. In *Seventh International AAAI Conference on Weblogs and Social Media (ICWSM 2013)*.
- [Cohen and Ruths 2013] Cohen, R., and Ruths, D. 2013. Classifying political orientation on Twitter: Its not easy! In *International AAAI Conference on Weblogs and Social Media*.
- [Conover et al. 2011] Conover, M.; Ratkiewicz, J.; Francisco, M.; Gonçalves, B.; Flammini, A.; and Menczer, F. 2011. Political polarization on Twitter. In *Proc. 5th International AAAI Conference on Weblogs and Social Media (ICWSM)*.
- [Garimella et al. 2016] Garimella, K.; De Francisci Morales, G.; Gionis, A.; and Mathioudakis, M. 2016. Quantifying controversy in social media. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining, WSDM '16*, 33–42. New York, NY, USA: ACM.
- [Garimella et al. 2017] Garimella, K.; Morales, G. D. F.; Gionis, A.; and Mathioudakis, M. 2017. Balancing opposing views to reduce controversy. In *Proceedings of the Tenth ACM International Conf. on Web Search and Data Mining, WSDM '17*. ACM.
- [Garimella, Weber, and Choudhury 2016] Garimella, K.; Weber, I.; and Choudhury, M. D. 2016. Quote RTs on twitter: usage of the new feature for political discourse. In *Proceedings of the 8th ACM Conference on Web Science (WebSci)*, 200–204.
- [Giatsoglou et al. 2015] Giatsoglou, M.; Chatzakou, D.; Shah, N.; Faloutsos, C.; and Vakali, A. 2015. Retweeting activity on twitter: Signs of deception. In *Advances in Knowledge Discovery and Data Mining - 19th Pacific-Asia Conference, PAKDD 2015, Ho Chi Minh City, Vietnam.*, 122–134.
- [Joshi et al. 2016] Joshi, A.; Tripathi, V.; Patel, K.; Bhattacharyya, P.; and Carman, M. J. 2016. Are word embedding-based features useful for sarcasm detection? In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, 1006–1011.
- [Kloumann and Kleinberg 2014] Kloumann, I. M., and Kleinberg, J. M. 2014. Community membership identification from small seed sets. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14*, 1366–1375. New York, NY, USA: ACM.
- [Kunegis et al. 2010] Kunegis, J.; Schmidt, S.; Lommatzsch, A.; Lerner, J.; Luca, E. W. D.; and Albayrak, S. 2010. Spectral analysis of signed graphs for clustering, prediction and visualization. In *Proc. SIAM Int. Conf. on Data Mining*, 559–570. SIAM.
- [Kwak et al. 2010] Kwak, H.; Lee, C.; Park, H.; and Moon, S. 2010. What is twitter, a social network or a news media? In *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, 591–600. New York, NY, USA: ACM.
- [Lanagan and Smeaton 2011] Lanagan, J., and Smeaton, A. F. 2011. Using twitter to detect and tag important events in live sports. *Artificial Intelligence* 542–545.
- [Lazer 2015] Lazer, D. 2015. The rise of the social algorithm. *Science* 348:1090–1091.
- [Leskovec, Huttenlocher, and Kleinberg 2010] Leskovec, J.; Huttenlocher, D.; and Kleinberg, J. 2010. Predicting positive and negative links in online social networks. In *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, 641–650. New York, NY, USA: ACM.
- [Liao, Wai-Tat, and Strohmaier 2016] Liao, Q.; Wai-Tat, F.; and Strohmaier, M. 2016. #snowden: Understanding biased introduced by behavioral differences of opinion groups on social media. In *Proceedings of the SIGCHI, CHI '16*. ACM.
- [Liu and Weber 2014] Liu, Z., and Weber, I. 2014. Predicting ideological friends and foes in twitter conflicts. In *23rd WWW*.
- [Livne et al. 2011] Livne, A.; Simmons, M. P.; Adar, E.; and Adamic, L. A. 2011. The party is over here: Structure and content in the 2010 election. In Adamic, L. A.; Baeza-Yates, R. A.; and Counts, S., eds., *ICWSM. The AAAI Press*.
- [Lo et al. 2011] Lo, D.; Surian, D.; Zhang, K.; and Lim, E.-P. 2011. Mining direct antagonistic communities in explicit trust networks. In *Proceedings of the 20th ACM International Conference on Information and Knowledge Management, CIKM '11*, 1013–1018. New York, NY, USA: ACM.
- [McGlone 2005] McGlone, M. S. 2005. Contextomy: the art of quoting out of context. *Media, Culture & Society* 27(4):511–522.
- [McPherson, Smith-Lovin, and Cook 2001] McPherson, M.; Smith-Lovin, L.; and Cook, J. M. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27(1):415–444.
- [Metaxas et al. 2015] Metaxas, P. T.; Mustafaraj, E.; Wong, K.; Zeng, L.; O'Keefe, M.; and Finn, S. 2015. What do retweets indicate? results from user survey and meta-review of research. In *Proceedings of the Ninth International Conference on Web and Social Media, ICWSM 2015, Oxford, UK*, 658–661.
- [Metaxas, Mustafaraj, and Gayo-Avello 2011] Metaxas, P. T.; Mustafaraj, E.; and Gayo-Avello, D. 2011. How (not) to predict elections. In *2011 IEEE Third International Conference on and 2011 IEEE Third International Conference on Social Computing (SocialCom), Boston, MA, USA, 2011*, 165–171. IEEE.
- [Mowbray 2010] Mowbray, M. 2010. The twittering machine. In Filipe, J., and Cordeiro, J., eds., *WEBIST (2)*, 299–304. INSTICC Press.
- [Mustafaraj and Metaxas 2011] Mustafaraj, E., and Metaxas, P. T. 2011. What edited retweets reveal about online political discourse. In *Analyzing Microtext*, volume WS-11-05 of *AAAI Workshops*. AAAI.
- [Rost et al. 2013] Rost, M.; Barkhuus, L.; Cramer, H.; and Brown, B. 2013. Representation and communication: Challenges in interpreting large social media datasets. In *Proceedings of the 2013 Conference on Computer Supported Cooperative Work, CSCW '13*, 357–362. New York, NY, USA: ACM.
- [Smith et al. 2014] Smith, M.; Rainie, L.; Shneiderman, B.; and Himelboim, I. 2014. Mapping twitter topic networks: From polarized crowds to community clusters. Pew Research Center. Last Accessed On 2017/01/05.
- [Tan et al. 2011] Tan, C.; Lee, L.; Tang, J.; Jiang, L.; Zhou, M.; and Li, P. 2011. User-level sentiment analysis incorporating social networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery*

- and data mining, KDD '11, 1397–1405. New York, NY, USA: ACM.
- [Tong, Faloutsos, and Pan 2008] Tong, H.; Faloutsos, C.; and Pan, J.-Y. 2008. Random walk with restart: fast solutions and applications. *Knowl. Inf. Syst.* 14(3):327–346.
- [Tufekci 2014] Tufekci, Z. 2014. Big questions for social media big data: Representativeness, validity and other methodological pitfalls. In *Proceedings of the 8th International Conference on Weblogs and Social Media, ICWSM, Ann Arbor, Michigan, USA*.
- [Vydiswaran et al. 2012] Vydiswaran, V. G. V.; Zhai, C.; Roth, D.; and Pirolli, P. 2012. Biastrust: teaching biased users about controversial topics. In wen Chen, X.; Lebanon, G.; Wang, H.; and Zaki, M. J., eds., *CIKM*, 1905–1909. ACM.
- [Wallace 2013] Wallace, B. 2013. Computational irony: A survey and new perspectives. *Artificial Intelligence Review* 1–17.
- [Weng et al. 2010] Weng, J.; Lim, E.-P.; Jiang, J.; and He, Q. 2010. Twiterrank: Finding topic-sensitive influential twitterers. In *Proceedings of the Third ACM International Conference on Web Search and Data Mining, WSDM '10*, 261–270. New York, NY, USA: ACM.
- [Wong et al. 2013] Wong, F. M. F.; Tan, C. W.; Sen, S.; and Chiang, M. 2013. Quantifying political leaning from tweets and retweets. In *Proceedings of the Seventh International Conference on Weblogs and Social Media, ICWSM 2013, Cambridge, Massachusetts, USA*.
- [Yang, Zhao, and Liu 2015] Yang, B.; Zhao, X.; and Liu, X. 2015. Bayesian approach to modeling and detecting communities in signed network. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, Austin, Texas, USA*.
- [Ye et al. 2013] Ye, J.; Cheng, H.; Zhu, Z.; and Chen, M. 2013. Predicting positive and negative links in signed social networks by transfer learning. In *Proceedings of the 22nd International Conference on World Wide Web, WWW '13*, 1477–1488.